

# BOG 11: Crosscut Portability Code Reuse Performance Portability

ASCR Workshop on Extreme Heterogeneity in HPC  
23-25 Jan 2018

# BOG 11 Contributors

Moderator(s): Anshu Dubey, Xiaoye Li, David Bernholdt

BOGists:

Rob Falgout

David Richards

Sam Williams

Scott Baden

Stan Tomov

Tiffany Mintz

Travis Humble

Wes Bethel

Christine Sweeney

Meifeng Lin

Line Pouchard

# BOG 11 Capability Targets for Extreme Heterogeneity

BOG 11 brainstorming and discussion of capabilities that will be needed in the 2025-2035 timeframe to make increasingly heterogeneous hardware technologies useful and productive for science applications.

# BOG 11 Targets for 2030

Target 1:

Target 2:

# BOG 11 Current Status

- Agreement that abstraction layers are needed for portability
  - Success stories are there, mostly groups that rolled their own
  - Ongoing efforts in Programming models, DSLs, runtimes, auto-tuning etc
  - Interplay between abstracting away the complexity (what we want) and what we have to deal with micromanagement of data placement, algorithms
- Currently gap between research and adoption
- Understanding that we need to have a broad coalition of CS researchers, application developers, and language standards to perhaps solve this problem.
  - We may repeat the history of GPUs if we don't solve this problem.

# BOG 11 Challenge Assessment (1/3)

Many topics overlap with other BOGs, specifically 5 and

- Software Design
  - Layers abstractions and their interplay
  - Separation of Concerns
- Barriers to adoption of abstractions
  - Relates back to how to use meaningfully in applications
- Provenance
  - Capturing execution details at runtime and post-hoc, analyzing data to get meaningful insights
  - Heisenberg problem (measurement affecting behavior)?
- Estimating costs
  - Of components within the applications, their interactions with machine components

# BOG 11 Challenge Assessment (2/3)

- Composability of components
  - At the level of using libraries
- Meta information about data or execution patterns useful to compilers and other tools
  - How to handle data locality in a meaningful way
- Measuring and understanding performance
  - Performance counters: how fine a granularity? Whether we capture the information we need? Can we actually get access to the counter info (privilege/engineering)?
  - Capturing the producer-consumer relationships between multiple accelerators (often different types)
- Understanding what portability implies with extreme heterogeneity
  - Just being able to run Vs using the resources in a useful way

# BOG 11 Challenge Assessment (3/3)

Challenges identified by other BOGs that apply here

- Mapping application heterogeneity to platform heterogeneity
- Building knowledge base of how and where abstraction can be adopted and how they interoperate
- API for interoperability at various interfaces
- Abstract machine model (at what level)



# BOG 11: list of key research challenges

Challenge 11.1 -- How to achieve separation of concerns

Challenge 11.2 -- How many layers of abstractions and where, and how do we express these abstractions and how do we incorporate them into applications

Challenge 11.3 -- How to create the coalition of stakeholders that can solve the problem of useful tools development and their adoption

Challenge 11.4 -- How to extract and understand performance, diagnostics, and cost of various components in ways that gives actionable insights.

Challenge 11.5 -- How to communicate meta-information about the application data and execution behavior to compilers and auto-tuning tools

# BOG 11 Possible Research Directions Summary

PRD 11.1 - Interdisciplinary research for software design that also informs the tools and abstractions design.

Possibly includes semantics constraining for languages in use

PRD 11.2 - direction 2 - Use of provenance to extract and understand information about application behaviour so that cost models can be built.

PRD 11.3 - direction 3 - Mechanism for communicating performance hints to compilers and mapping tools (distinct from constructs in languages)

# PRD 11.n : Short title of possible research direction

- One paragraph description (3 sentence/bullet)
- Research challenges
  - Metrics for progress
- Potential research approaches and research directions
- How and when will success impact technology?